arXiv:1612.08879v1 [cs.CV] 28 Dec 2016

Deep Unsupervised Representation Learning for Remote Sensing Images

DaoYu Lin
lindaoyu15@mails.ucas.ac.cn

Institute of Electronics, Chinese Academy of Sciences

Abstract Scene classification plays a key role in interpreting the remotely sensed high-resolution images. With the development of deep learning, supervised learning in classification of Remote Sensing with convolutional networks (CNNs) has been frequently adopted. However, researchers paid less attention to unsupervised learning in remote sensing with CNNs. In order to filling the gap, this paper proposes a set of CNNs called **M**ultiple l**A**ye**R** fea**T**ure m**A**tching(MARTA) generative adversarial networks (GANs) to learn representation using only unlabeled data. There will be two models of MARTA GANs involved: (1) a generative model G that captures the data distribution and provides more training data; (2) a discriminative model D that estimates the possibility that a sample came from the training data rather than G and in this way a well-formed representation of dataset can be learned. Therefore, MARTA GANs obtain the state-of-the-art results which outperform the results got from UC-Merced Land-use dataset and Brazilian Coffee Scenes dataset.

Keywords: Unsupervised representation learning, generative adversarial networks, scene classification

1 Introduction

With the improvement of satellite imaging techniques, an ever-growing number of high-resolution satellite images that are provided by special satellite sensors become available. It is of great necessity to interpret such massive image repositories in an automatic and accurate way. In the past decades, remote sensing scene classification has become a heated topic and a fundamental method for application among various applications such as land-resource management and urban planning. CNNs (Penatti et al., 2015) has achieved greater success in classification of remote sensing images with the massive data. However, the cost for labeled remote sensing images is very high. Therefore, we explore an unsupervised approach based on adversarial training inspired by the generative adversarial networks (GANs) approach (Goodfellow et al., 2014). In order to make good image representations, we proposed a *feature matching layer* pooling from multiple layers to fuse coarse, semantic and local, appearance information. The overall model is depicted in the following Fig. 1.

The contributions of this paper are the following:

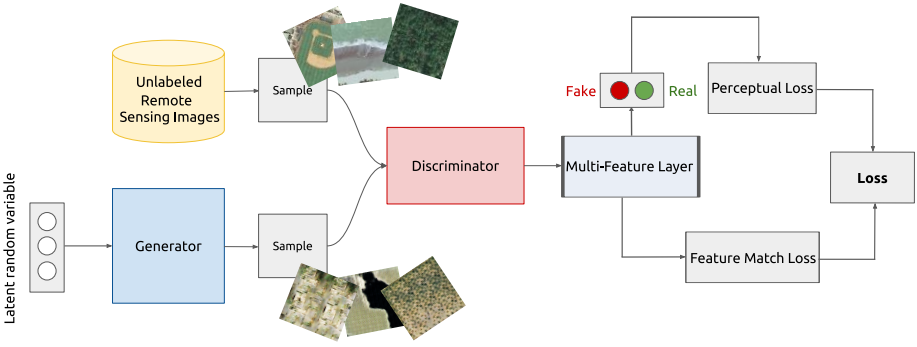

Figure 1: Overview of the proposed approach. The discriminator learns to classify between real and synthesized images, while the generator learns to fool the discriminator. Adversarial net takes label map as input and produces class label (1=real, or 0=synthetic).

1. It is the first time to apply the adversarial training in the field of remote sensing images.
2. This method obtains the state-of-the-art classification results in UC-Merced Land-use dataset and Brazilian Coffee Scenes dataset.
3. The generator G gets the unprecedented resolution (256×256) optical remote-sensing images and green-red-infrared remote-sensing images for the first time.

The outline of the paper will be discussed. Section 2 develops the related work on unsupervised representation learning and adversarial training approaches. After presenting our adversarial training approach and MARTA GANs architecture in Section 3, experimental results will be discussed in Section 4. Finally, a conclusion will be made in Section 5.

2 Related Work

Unsupervised representation learning. The purpose of unsupervised representation learning is to use unlabeled data to learn a representation that exposes important semantic features as easily decodable factors. Unsupervised representation learning in a well-studied problem in the field of general computer vision. A standard approach to unsupervised representation learning is to do clustering based on the data, and leverage the clusters for improved classification scores. In the context of images, one can do hierarchical clustering of image patches (Coates & Ng, 2012) to learn powerful image representations. Another popular method is to train auto-encoders, stacking layers of denoising autoencoders (Vincent et al., 2010) which are trained locally to denoise corrupted versions of their inputs , separating the what and where components of the code (Zhao et al., 2015), ladder structures (Rasmus et al., 2015) that encode an image into a compact

code, and decode the code to reconstruct the image as accurately as possible. A lot of promising recent work originates from the Skip-gram model (Mikolov et al., 2013), which inspired the skip-thought vectors (Kingma & Ba, 2015) and several techniques for unsupervised feature learning of images (Doersch et al., 2015).

Generative adversarial networks. The goal of adversarial approach is to learn deep generative models. The generator network directly produces samples $x = G(z; \theta_g)$. However, the discriminator network endeavors to make a distinction between samples drawn from the training data and samples drawn from the generator. The discriminator produces a probability value given by $D(x; \theta_g)$, demonstrating the probability that x is a real training example instead of a fake sample drawn from the model. In follow-up work, a deep convolutional GAN (DCGAN) (Radford et al., 2015) performed very well in image synthesis tasks, and showed that it's latent representation space captures important factors of variation. A great variety of new architectural features and training procedures (Salimans et al., 2016), including feature matching and minibatch discrimination, which can produce very clean and sharp images and learn codes that contain valuable information about these textures. InfoGAN (Chen et al., 2016), an information-theoretic extension to the Generative Adversarial Network that is able to learn disentangled representations in a completely unsupervised manner.

3 Method

3.1 Adversarial learning

The adversarial framework is most straightforward to apply when the models are both multilayer perceptrons. To learn the generators distribution p_g over data x, we define a prior on input noise variables $p_z(z)$, then represent a mapping to data space as $G(z; \theta_g)$, where G is a differentiable function represented by a multilayer perceptron with parameters θ_g . We also define a second multilayer perceptron $D(x; \theta_d)$ that outputs a single scalar. $D(x)$ represents the probability that x came from the data rather than p_g. We train D to maximize the probability of assigning the correct label to both training examples and samples from G, simultaneously train G to minimize $\log(1 - D(G(z)))$. We called this for perceptual loss, it encourages the reconstructed image to be similar to the samples drawn from the training set.

$$\ell_{perceptual} = \mathbb{E}_{z \sim p_z(z)}[\log(1 - D(G(z)))] \tag{1}$$

In other words, D and G play the following two-player minimax game with value function $V(D, G)$:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{\text{data}}(x)} \log D(x) + \mathbb{E}_{z \sim p_z(z)}[\log(1 - D(G(z)))] \tag{2}$$

3.2 Feature matching

Feature matching addresses the instability of GANs by specifying a new objective for the generator that prevents it from over training on the current discriminator.

Instead of directly maximizing the output of the discriminator, this objective requires the generator to generate data that matches the statistics of the real data, where we use the discriminator only to specify the statistics that we think are worth matching. Specifically, we train the generator to match the expected value of the features on the *feature matching layer* of the discriminator. This is a natural choice of statistics for the generator to match, since by training the discriminator we ask it to find those features that are most discriminative of real data versus data generated by the current model.

Letting $f(x)$ denote activations on an intermediate layer of the discriminator, the objective for the generator is defined as:

$$\ell_{feature_match} = ||\mathbb{E}_{x \sim p_{\text{data}}(x)} f(x) + \mathbb{E}_{z \sim p_z(z)} f(G(z))||_2^2 \tag{3}$$

The discriminator, and hence $f(x)$, are trained in the usual way. As with regular GAN training, the objective has a fixed point where G exactly matches the distribution of training data. So, our final objective (combined Eqn. 1 and Eqn. 3) for discrimininator is

$$\ell_{final} = \ell_{perceptual} + \lambda \ell_{feature_match} \tag{4}$$

3.3 Network architectures

After extensive model exploration we identified a family of architectures that resulted in stable training across a range of datasets and allowed for training higher resolution remote sensing images and deeper generative models. The details of generator and discriminator in MARTA GANs are in the below.

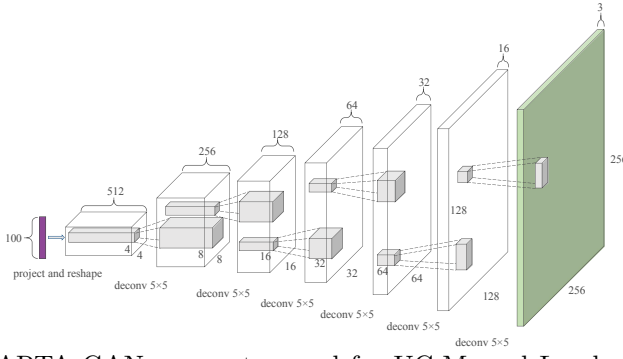

Figure 2: MARTA GANs generator used for UC-Merced Land-use dataset. A 100 dimensional uniform distribution $p_z(z)$ is projected to a small spatial extent convolutional representation with many feature maps. A series of six fractionally-strided convolutions then convert this high level representation into a 256×256 pixel image.

The first layer of the generator, which regards a uniform noise distribution z as input, could be called fully connected as it is just a matrix multiplication,

but the result is reshaped into a 4-dimensional tensor and used as the start of the convolution stack. We use fractional-strided convolutions (also called deconvolutions) (Dosovitskiy et al., 2015) in our generator, allowing it to learn its own spatial upsampling, this paper uses six layers fractional-strided convolutions, upsampled 4×4 feature map to 256×256 remote sensing image. The ReLU activation (Nair & Hinton, 2010) is used in the generator with the exception of the output layer which uses the tanh function. See Fig. 2 for a visualization of an example model architecture.

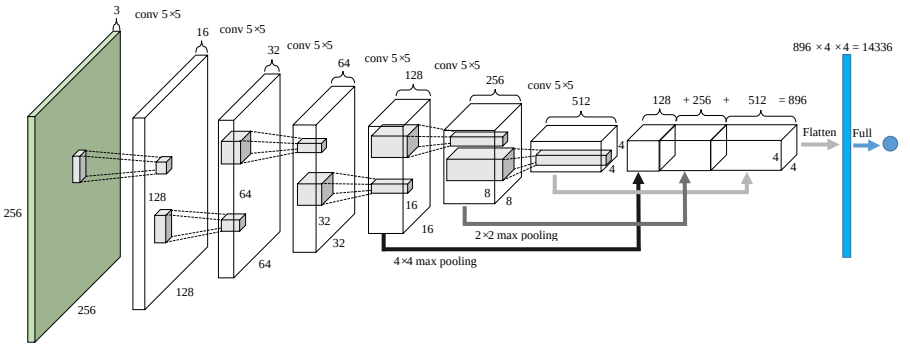

Figure 3: MARTA GANs discriminator used for UC-Merced Land-use dataset. The input includes real or synthetic images. A series of six strided convolutions then maxpooling last three layers representation in order to produce a 4×4 spatial grid and flatted it as representation of the input image.

For the discriminator, the first layer is input images which includes both real image and synthetic image. We use strided convolutions in our discriminator which allows it to learn its own spatial downsampling. What's more, we developed a new layer(multiple feature layer) which maxpooling last three layers representation in order to produce a 4×4 spatial grid and concatenated them to a new layer. We use LeakyReLU activation in the discriminator for all layers. The last convolution layer is flattened and then fed into a single sigmoid output. See Fig. 3 for a visualization of an example model architecture.

We used batch normalization Ioffe & Szegedy (2015) for both generator and discriminator which stabilizes learning by normalizing the input to each unit to have zero mean and unit variance. This helps to deal with training problems that arise due to poor initialization and helps gradient flow in deeper models.

4 Experiments

We trained MARTA GANs on two datasets, UC Merced Land Use Dataset and Brazilian Coffee Scenes Dataset. We implemented MARTA GANs by TensorLayer [1], a deep learning and reinforcement learning library extended from Google

[1] `http://tensorlayer.readthedocs.io/en/latest/`

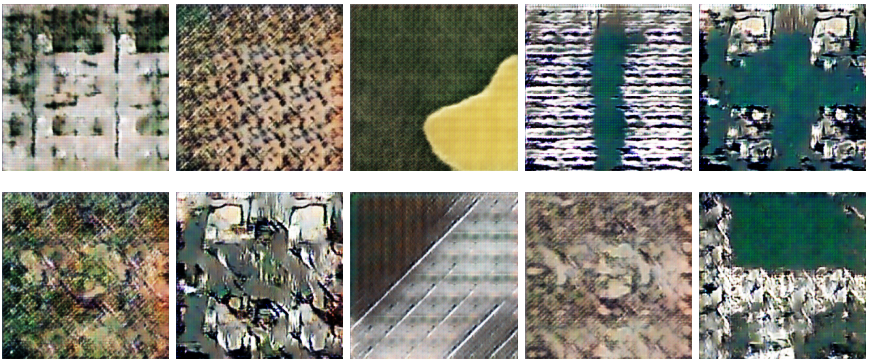

Figure 4: Exemplary images produced by a generator trained on UC-Merced using the ℓ_{fianl} (Eqn. 4) objective. Generator sampled with unprecedented resolution (256×256) Remote Sensing Images.

TensorFlow (Abadi et al., 2016). Details on the usage of each of these datasets are given below. In all cases we used a moderate data augmentation, through flip horizontally and vertically, rotation 90 degrees to increase the effective training set size. No pre-processing was applied to training images besides scaling to the range of the tanh activation function [-1, 1]. All models were trained with mini-batch stochastic gradient descent (SGD) with a mini-batch size of 64. All weights were initialized from a zero-centered Normal distribution with standard deviation 0.02. In the LeakyReLU, the slope of the leak was set to 0.2 in all models. While previous GAN work has used momentum to accelerate training, we used the Adam optimizer (Kingma & Ba, 2015) with tuned hyperparameters and the learning rate is 0.0002. Additionally, we set the momentum term $\beta_1 = 0.5$ to help stabilize training. And we set BatchNormLayer decay factor 0.9 for Exponential Moving Average.

One common technique for evaluating the quality of unsupervised representation learning algorithms is to apply them as a feature extractor on supervised datasets and to evaluate the performance of linear models fitted on top of these features. Our aim is to evaluate whether MARTA GANs can learn good representations or not.

4.1 UC Merced

The dataset (Yang & Newsam, 2010) consists of images of 21 land-use classes selected from aerial optical images of the US Geological Survey, taken over various regions of the United States. 100 images measuring 256×256 pixels were manually selected for 21 land use classes, 100 for each class. More information are available in the original paper (Yang & Newsam, 2010).

To evaluate the quality of the representations learned by MARTA GANs for supervised tasks, we trained on UC-Merced maxpooling the last three layers representation to produce a 4×4 spatial grid. These features are then flattened

Table 1: Classification accurary (%) of representation extracted by DCGANs and MARTA GANS on the UC-Merced dataset. Best result in bold.

Architectures	Design	Accuracy
DCGANs	Without data augmentation	80.36
	With data augmentation	87.01
MARTA GANs	Without data augmentation	85.37
	With data augmentation, $\ell_{perceptual}$ only	93.57
	With data augmentation, use ℓ_{final}	**95.00**

and concatenated to a 14336 dimensional vector and a regularized linear L2-SVM classifier is trained on top of them. This achieves 95.00% accuracy, outperforming all unsupervised representation Learning methods on this dataset.

In Table 1 we report synthetic results for the two GANs architectures and weather used data augmentation. The first observation is that with data augmentation approach, as expected, provides the best results with both DCGANs and MARTA GANs, reaching an overall accuracy of 87.01% and 95.00% respectively. This is about 7% and 10% better than the without data augmentation, respectively. The data augmentation is an effective way to reduce overfitting when training a large deep network, which generates more training image samples by rotation and flipping these patches from original images. We also evaluate the performance of the generator network for two losses $\ell_{perceptual}$ (Eqn. 1)and ℓ_{final} (Eqn. 3), combined this two losses (in Eqn. 4 we use $\lambda = 1.0$) achieved the best performance, and generated remote sense examples provided in Fig. 4.

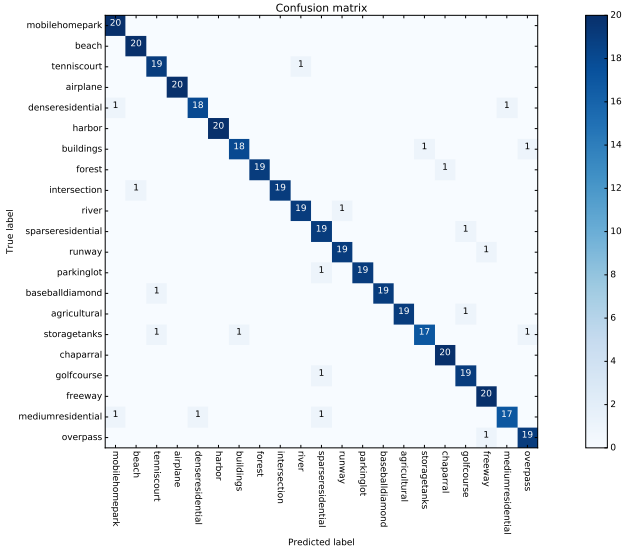

Figure 5: Confusion matrices on the UC-Merced dataset using features extracted from MARTA GAN.

Fig. 5 shows the confusion matrices of the representation features of discriminator learned. The features extraction form multiple feature layer in discriminator result in 100% accuracy for most of the scene categories. Comparing the presence with other very close classes, like denseresidential, building, medium residential, get about 88.33% accuracy.

In addition, we visualize the image global representations encoded via MARTA GANs features for the UC-Merced dataset. Here, we compute features for all image scenes in the dataset first, and then use the t-SNE algorithm (Maaten & Hinton, 2008) to embed the high-dimensional image features on a 2-D space. We show these 2-D embedding points with different colors and shapes corresponding to their actual scene categories. The final visualization results are shown in Fig. 6. This observation show that features extracted from matri-feature layers are high-level features that contain abstract semantic information.

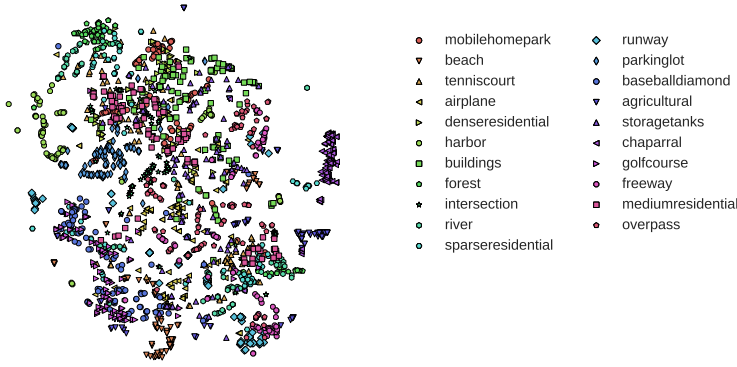

Figure 6: 2-D feature visualization of image global representations of the UC-Merced dataset.

Table 2: Classification accuracy (%) of reference and proposed methods on the UC-Merced dataset. Best result in bold.

Method	Accuracy
HMFF (Shao et al., 2013)	92.38
VLAT (Negrel et al., 2014)	94.30
UFL-SC (Hu et al., 2015)	90.26
PSR (Chen & Tian, 2015)	89.10
MCMI-based (Ren et al., 2015)	88.20
MARTA GANs(proposed)	**95.00**

Several approaches have been proposed recently for remote sensing scene classification, and most of them have been tested on the UC-Merced dataset. Therefore, there is a lot of data available for a solid comparison with the state-of-the-art. In Table 2 we report the overall accuracies for all these comparable methods, as they appear in the original papers, together with the accuracy of

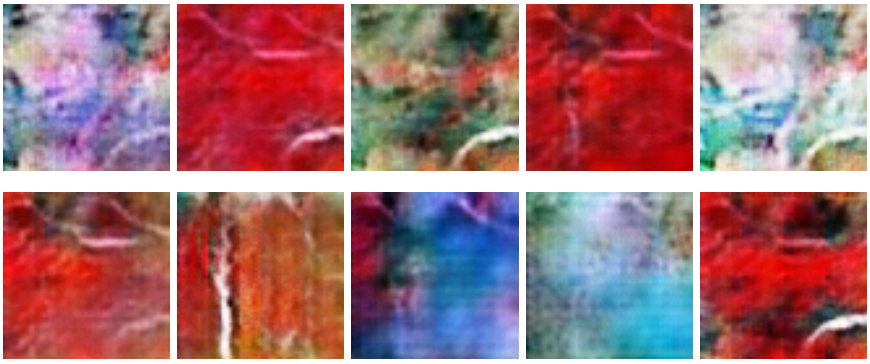

Figure 7: Exemplary images produced by a generator trained on Brazilian Coffee dataset using the ℓ_{fianl} (Eqn. 4) objective. Is the first time generator sampled with green-red-infrared remote sensing images.

our best MARTA GANs method. The proposed method guarantees a better performance gain w.r.t. to all references.

4.2 Brazilian Coffee Scenes

This dataset (Penatti et al., 2015) is a composition of scenes taken by SPOT sensor in 2005 over four counties in the State of Minas Gerais, Brazil: Arceburgo, Guaranésia, Guaxupé and Monte Santo. The whole image set of each county was partitioned into multiple tiles of 64×64 pixels. For this dataset, it was considered only the green, red, and near-infrared bands, which are the most useful and representative ones for discriminating vegetation areas. The creation of the dataset is performed as follows: tiles with at least 85% of coffee pixels were assigned to the coffee class; tiles with less than 10% of coffee pixels were assigned to the non-coffee class, it has 1438 tiles of coffee and 1438 tiles of uncoffee. Fig. 7 shows some examples produced by a generator trained on this dataset . Overall, this dataset is very different from UC-Merced. The images are not optical (green-red-infrared instead of red-green-blue).

Table 3: Classification accurary (%) of representation extracted by DCGANs and MARTA GANS on the Coffee Scenes dataset. Best result in bold.

Architectures	Design	Accuracy
DCGANs	Without data augmentation	85.36
	With data augmentation	85.01
MARTA GANs	Without data augmentation	87.69
	With data augmentation, $\ell_{perceptual}$ only	87.73
	With data augmentation, use ℓ_{final}	**88.36**

Table 3 shows the results obtained with the proposed techniques. The most notable difference w.r.t. the UC-Merced case is the data augmentation doesnt

improve the accuracy conspicuously. In fact, this just a 2-class problem, and it has enough data to train the network. In general, results are significantly worse than with UC-Merced, despite the 2-class vs. 21-class problem. This is a rather challenging dataset, due to a large intra-class variability caused by different crop management techniques, different plant ages and/or spectral distortions and shadows. In any case, this result is better than the other unsupervised top result, obtained with BIC (Penatti et al., 2015) (Border-Interior Pixel Classification) a simple color descriptor, see Table 4.

Table 4: Classification accuracy (%) of reference and proposed methods on the Coffee Scenes dataset. Best result in bold.

Method	Accuracy
BIC (Penatti et al., 2015)	87.00
BOVM (Penatti et al., 2015)	80.50
MARTA GANs(proposed)	**88.36**

5 Discussion

This paper introduces a representation learning algorithm called **M**ultiple l**A**ye**R** fea**T**ure m**A**tching(MARTA) generative adversarial networks (GANs). In contrast to previous approaches, which require supervision, MARTA GANs is completely unsupervised and learns interpretable representations on challenging remote sensing datasets. In addition, MARTA GANs introduce a new multiple feature matching layer to learn multi-scale spatial information for high-resolution remote sensing. Other possible extensions to this paper include: producing high quality samples of remote sensing images by the generatorlearning hierarchical latent representations, classifying remote sensing image with semi-supervised to get a better accuracy.

Bibliography

Abadi, Martın, Agarwal, Ashish, Barham, Paul, Brevdo, Eugene, Chen, Zhifeng, Citro, Craig, Corrado, Greg S, Davis, Andy, Dean, Jeffrey, Devin, Matthieu, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.

Chen, Shizhi and Tian, YingLi. Pyramid of spatial relatons for scene-level land use classification. *IEEE Transactions on Geoscience and Remote Sensing*, 53 (4):1947–1957, apr 2015. doi: 10.1109/tgrs.2014.2351395.

Chen, Xi, Duan, Yan, Houthooft, Rein, Schulman, John, Sutskever, Ilya, and Abbeel, Pieter. InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets. *arXiv*, 2016. URL `http://arxiv.org/abs/1606.03657`.

Coates, Adam and Ng, Andrew Y. Learning feature representations with K-means. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 7700 LECTU: 561–580, 2012. ISSN 03029743. doi: 10.1007/978-3-642-35289-8-30.

Doersch, Carl, Gupta, Abhinav, and Efros, Alexei a. Unsupervised Visual Representation Learning by Context Prediction. *arXiv*, pp. 1422–1430, 2015. ISSN 978-1-4673-8391-2. doi: 10.1109/ICCV.2015.167. URL `http://arxiv.org/abs/1505.05192`.

Dosovitskiy, Alexey, Tobias Springenberg, Jost, and Brox, Thomas. Learning to generate chairs with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1538–1546, 2015.

Goodfellow, Ian J., Pouget-Abadie, Jean, Mirza, Mehdi, Xu, Bing, Warde-Farley, David, Ozair, Sherjil, Courville, Aaron, and Bengio, Yoshua. Generative Adversarial Networks. *arXiv preprint arXiv: ...*, pp. 1–9, jun 2014. ISSN 10495258. URL `http://arxiv.org/abs/1406.2661`.

Hu, F, Xia, G, Wang, Z, Huang, X, Zhang, L, and Sun, H. Unsupervised Feature Learning Via Spectral Clustering of Multidimensional Patches for Remotely Sensed Scene Classification. *Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of*, 8(5):2015–2030, 2015. ISSN 1939-1404. doi: 10.1109/JSTARS.2015.2444405.

Ioffe, Sergey and Szegedy, Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

Kingma, Diederik P. and Ba, Jimmy Lei. Adam: a Method for Stochastic Optimization. *International Conference on Learning Representations 2015*, pp. 1–15, 2015.

Maaten, Laurens van der and Hinton, Geoffrey. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.

Mikolov, Tomas, Chen, Kai, Corrado, Greg, and Dean, Jeffrey. Efficient Estimation of Word Representations in Vector Space. *Arxiv*, pp. 1–12, 2013. ISSN

15324435. doi: 10.1162/153244303322533223. URL http://arxiv.org/abs/1301.3781.

Nair, Vinod and Hinton, Geoffrey E. Rectified Linear Units Improve Restricted Boltzmann Machines. *Proceedings of the 27th International Conference on Machine Learning*, (3):807–814, 2010. doi: 10.1.1.165.6419.

Negrel, Romain, Picard, David, and Gosselin, Philippe Henri. Evaluation of second-order visual features for land-use classification. *Proceedings - International Workshop on Content-Based Multimedia Indexing*, 2014. ISSN 19493991. doi: 10.1109/CBMI.2014.6849835.

Penatti, Otávio AB, Nogueira, Keiller, and dos Santos, Jefersson A. Do deep features generalize from everyday objects to remote sensing and aerial scenes domains? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 44–51, 2015.

Radford, Alec, Metz, Luke, and Chintala, Soumith. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv*, pp. 1–15, 2015. ISSN 0004-6361. doi: 10.1051/0004-6361/201527329. URL http://arxiv.org/abs/1511.06434.

Rasmus, Antti, Valpola, Harri, and Berglund, Mathias. Semi-Supervised Learning with Ladder Network. *arXiv*, pp. 1–17, 2015. ISSN 10495258. doi: 10.1017/CBO9781107415324.004.

Ren, Jianfeng, Jiang, Xudong, and Yuan, Junsong. Learning LBP structure by maximizing the conditional mutual information. *Pattern Recognition*, 48(10): 3180–3190, 2015. ISSN 00313203. doi: 10.1016/j.patcog.2015.02.001.

Salimans, Tim, Goodfellow, Ian, Zaremba, Wojciech, Cheung, Vicki, Radford, Alec, and Chen, Xi. Improved Techniques for Training GANs. *Nips*, pp. 1–10, 2016. doi: arXiv:1504.01391.

Shao, Wen, Yang, Wen, Xia, Gui Song, and Liu, Gang. A hierarchical scheme of multiple feature fusion for high-resolution satellite scene categorization. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 7963 LNCS:324–333, 2013. ISSN 03029743. doi: 10.1007/978-3-642-39402-7_33.

Vincent, Pascal, Larochelle, Hugo, Lajoie, Isabelle, Bengio, Yoshua, and Manzagol, Pierre-Antoine. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11(Dec):3371–3408, 2010.

Yang, Yi and Newsam, Shawn. Bag-of-visual-words and spatial extensions for land-use classification. In *Proceedings of the 18th SIGSPATIAL international conference on advances in geographic information systems*, pp. 270–279. ACM, 2010.

Zhao, Junbo, Mathieu, Michael, Goroshin, Ross, and LeCun, Yann. Stacked What-Where Auto-encoders. *arXiv preprint arXiv:1506.02351*, 1(i):1–12, 2015. URL http://arxiv.org/abs/1506.02351.